



0172.38689X00
NC 1718 US

SUBSTITUTE SPECIFICATON

METHOD FOR SCHEDULING PACKETIZED DATA TRAFFIC

BACKGROUND OF THE INVENTION

Field of the Invention

[0001] This invention relates to a method for scheduling data streams and, more particularly, to a method for scheduling a packet based data stream using a slack value calculation.

RECEIVED

JUN 23 2004

Description of the Prior Art

Technology Center 2600

[0002] In many networking arrangements, it is necessary for a plurality of data streams to be combined to share a limited number of channels or even a single line. This can happen, for example, in a wireless type network where a number of units are routed to a base station which is further connected through a single channel. Thus, the various streams going to the individual units must be handled through a single channel. Another similar situation is where various kinds of data streams, such as voice data, real time video data, email and other data are all handled through an Internet protocol network. These various different kinds of data must be combined into a limited number of channels or even a single channel for transmission.

[0003] Whenever such data streams are merged, it is necessary to have some protocol for selecting the order in which they are placed in the channel. While the simplest solution may be a first come, first serve arrangement, this may not be the most effective since some data streams are more time sensitive than others. For example, voice signals cannot be delayed very much at all, whereas email messages may be delayed by a substantial amount. Accordingly, a number of protocols have been sought to provide fair and optimum criteria for multiple users so that delay is reduced, invalid data is minimized and data throughput is maximized.

[0004] One attempt at such a scheduling method is referred to as deficit round robin where a fairness level is achieved by using a deficit counter and a quantum of service for each user flow which decides how long the flow should constantly be served before moving onto the next data flow. The maximum delay for revisiting a user is governed by the round duration in the scheduler and depends on the packet lengths and the number of flows in the system. However, this method is inefficient when lower bound delay requirements must be satisfied. This is because the packet may be delayed by a full round duration and since the maximum delay is governed by the round duration, it would be impossible to provide different delay bounds to different flows, thereby resulting in a high drop ratio of packets, that is, real time packets that exceed their delay requirements.

[0005] In a method called the weighted fair queuing, the delay in a user packet flow is decreased by increasing the allocated service rate and in another method referred to as earliest due date, each flow is served using a deadline base

strategy where the user with the packet of earliest deadline waiting to be scheduled is selected first. In these two systems, the transmission of a scheduled packet must be completed before scheduling another packet. Therefore, the delay guarantees of a packet depends on the length of another packet in a different flow sharing the same channel. Thus, a new short packet arriving in the system could time out while waiting for another packet of lower sensitivity to finish transmission. This leads to lower system throughput. Another drawback of weighted fair queuing scheduling is that the number of bits served in a scheduling round is proportional to the rate allocated to the flow. To reduce the delay for a flow, its allocated rate must be set-up to a higher volume before starting service of the flow. Given that the rate is fixed throughout service of the flow, the coupling between the rate allocation and the delay may lead to inefficient resource utilization. While a low value does not provide enough quality of service, a very large rate allocation leads to a waste of bandwidth. This is due to the rate fluctuations in a real time variable bit rate traffic.

SUMMARY OF THE INVENTION

[0006] Accordingly, the present invention provides a method for scheduling data flows from concurrent data streams.

[0007] The present invention also provides a method for scheduling data streams by splitting packets into data segments.

[0008] This invention further provides a method for scheduling data segments from different data streams using a slack measure for a scheduling decision.

[0009] The present invention still further provides the use of a slack measurement for data segment scheduling decisions where the slack value is based on its deadline and estimated transmission time.

[0010] The present invention further provides an apparatus for segmenting incoming data traffic packets into data segments, for assigning slack time to each data segment and for scheduling the transmission of the packets based on the determined slack time.

[0011] Briefly, the invention is achieved by providing a method for splitting a data packet into data segments which can be scheduled independently for transmission and using a slack measure as an input to the scheduling decision. The slack measurement provides a measure of how much cumulative time a packet can tolerate to wait and still meet its requirements.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] A more complete appreciation of the invention and many of the attendant advantages thereof will be readily obtained as the same becomes better understood by reference to the following detailed description when considered in connection with the accompanying drawings, wherein:

[0013] Fig. 1 is a graph showing the measurement of the slack time;

[0014] Fig. 2 is a block diagram showing an apparatus for scheduling data traffic according to the present invention;

[0015] Fig. 3 is a flow chart showing the steps of the method for scheduling transmission according to the present invention;

[0016] Fig. 4 is a flow chart indicating the steps involved in a variation of the method of Figure 3; and

[0017] Fig. 5 is a block diagram showing a second apparatus for scheduling data traffic according to the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0018] The present method is based on two steps. First, a data packet is split into data segments which can independently be scheduled for transmission. Secondly, a slack measure is calculated for each packet based on its quality of service requirements that could be assigned to the data segments to provide a measure of how much cumulative waiting time the packet can tolerate and still meet its necessary requirements. By splitting the packet into shorter data segments, it is possible to move parts of the packet independently which gives the scheduler more freedom to intersperse parts of the packet and to thus keep the flow more even. Thus, unlike other proposed methods, the present slack measure method supports the variation in delay bounds between different flows. This is due to the ability to monitor the requirements of each packet independently and, therefore, provide a dynamic priority allocation at the packet basis. Thus, it provides scheduling at the data segment level. This provides the scheduler with the capability of switching service between multiple packets. Thus, packets which are sensitive to delay do not have to wait until another packet transmission is complete if that packet can afford extra delay. This leads to an increase in the system throughput.

[0019] The present method also allows for possibilities to look ahead and determine if a packet will exceed its requirements and, therefore, drop the packet completely before continuing or starting its service. This is useful in congestion control and for efficient utilization of the bandwidth. This is achieved by eliminating the allocation of bandwidths for packets that are not successfully extracted at the user end and using the bandwidth for other packets that are known to be successfully generated at the user end. This improves the quality of service as seen by the user and increases bandwidth utilization efficiency.

[0020] The various data packets are split into data segments for scheduling and transmission. In packet cellular systems, the data segments correspond to the radio link control or multiple access control blocks. A data segment is transmitted individually over the transmission channel when a transmission opportunity is granted. Thus, in time division multiple access systems, a transmission opportunity is a time slot. In a wideband code division multiple access system, it is the utilization of the unique Walsh code in a radio frame. In this system, the radio frame is shared by multiple users using different Walsh codes.

[0021] Once the packet is split into data segments, it is necessary for the scheduler to schedule the data segments. The segments are organized into a schedule in order to use the transmission opportunities which are available. Thus, for example, in a wireless link, where concurrent users are involved, the scheduler would schedule concurrent user traffic flows waiting for service on both the uplink and the downlink. However, the scheduling must be accomplished so

that all user flows are served in a fair and optimum fashion that will meet the traffic constraints and maximize the system throughput.

[0022] In order to determine the order in which the data segments from different flows should be arranged, the present method assigns a slack value for each packet. This value is then similarly assigned to each segment from that packet. The slack value is calculated based on the deadline by which the packet must complete its transmission and also the estimated minimum transmission time. This is graphically shown in Fig. 1 where time is the variable in the horizontal direction starting at point O. The time necessary to transmit a packet, assuming no delays, is shown in the cross-hatched area extending from O to A. If the time deadline for transmitting the packet is given at point B, the slack time is the time between the transmission time and the deadline or as indicated in the Figure, AB. The slack value is a number which corresponds to this amount of time. Thus, the slack value is directly related to the amount of time that the segments from the packet can wait before transmission is started. This value can be measured in terms of the number of transmission opportunities that can be missed during the transmission of all the data segments of the packet.

[0023] After the packet is segmented into data segments and the slack value is assigned to each of the segments, the segments wait for their turn. One manner of using the slack value is to give lower slack value segments priority over higher values. Other methods may also be used by incorporating other criteria along with the slack value to determine priority. Every time the packet is not included in a transmission opportunity, its slack value is decreased. In terms of

Fig. 1, the deadline B becomes closer to the current time O but the transmission time of the remaining data segments of the packet does not change so that the slack time AB becomes smaller. This indicates that the amount of time it can wait before transmission decreases. When a part of the packet is included in a transmission, the slack value does not change. This is because the amount of transmission time needed will be shortened by the same amount of time that the deadline is shortened. Thus, in Figure 1, while B will get closer to O, A will also be closer to O by the same amount so that the slack time AB remains the same. When a segment has a slack value of zero, this indicates that it must be serviced at every upcoming transmission opportunity in order to meet its deadline requirement. Since the data segments of a packet all have the same value, these data segments all have a slack value of zero at the same time and thus they will all be transmitted at every transmission opportunity.

[0024] Given the real time nature and variable bit rate characteristics of the traffic (such as real time video and streaming video) the present slack method dynamically organizes the transmission order of the packets in the system in order to meet quality of service requirements involving delay and jitter.

[0025] This scheduling method is integrated with other systems and methods in the transmission device which also operate to control the transmission of the data. However, these other systems are not discussed herein since they do not effect the particular operation of the scheduling method.

[0026] Fig. 2 shows an apparatus 10 for accomplishing this process. A plurality of data streams, labeled input user data traffic and indicated by the

arrows on the left hand side of the figure, are input into the system. They first enter a segmenter 12 which divides the packets into a series of data segments. The segments are stored in queues 14. The packets are also assigned a slack time value as discussed above, by the slack time assigner 16. This information is stored, then used by the scheduler 18 which selects the data segments which are to be transmitted next. The scheduler outputs the selected data segments and sends them to transmitter 20 for transmission. It is possible for a single user to have more than one active data stream. Thus, more than one queue could hold data for the same user.

[0027] Fig. 3 is a flowchart showing the basic steps of the method described above. That is, in step 30, the incoming data traffic is segmented in order to form data segments. The slack time value is calculated in step 31 and then assigned to the packets in step 32. This value corresponds to the measurement of the slack time as discussed above in regard to Figure 1. Based on the slack time, the schedule of data segments is determined in step 34. For segments which are not going to be transmitted immediately, the slack time value is recalculated in step 36 and the new slack time is inserted in step 32. Once the particular data segment is selected, it is then transmitted in step 38.

[0028] Fig. 4 shows a procedure which can be included in step 34 to determine if a packet does not have any chance of being sent in time. In such a situation, it is better to not bother to waste time sending part of it since it will be useless at the other end anyway and since the time can better be spent on generating other segments. Accordingly, when a segment is being considered for the schedule in

step 34 in Fig. 3, instead the arrangement in Fig. 4 can be used. That is, first the packet is examined to determine if it will exceed the amount of time that it has available and thus not be suitable for sending. If it does not exceed these requirements, it then is scheduled for transmission in the normal course of the method as described above. However, if it does exceed these requirements then the entire packet is deleted as indicated in step 42.

[0029] Fig. 5 shows an alternative arrangement 110 to the apparatus of Fig. 2.

Input user data traffic is still indicated by arrows on the left hand side of the figure. The data traffic first enters a segment calculator which calculates the slack time for the entire packet. The entire packet is then placed in queues 114.

Segmenter 112 operates on the packets as they reach the front of the queue and after the scheduler 118 selects it for transmission. The segmenter 112 will generate only one data segment from the packet based on the data segment length available from the transmitter for the current transmission opportunity. Thus, the data segment limit length can vary at any time. The transmitter is aware of this change and will send information regarding the new data segment length value to the segmenter and the segment calculator. Thus, in this arrangement the slack value is assigned to the entire packet and not a single segment. Thus, all segments of the same packet have the same slack value. Any recalculating of the slack value for one segment requires the same value for all of their data segments in the same packet.

[0030] The slack process described above can be implemented at any node in the network where scheduling is required so that different delay guarantees or packets of different traffic flows can be accomplished. It can also be used in scheduling non-real time traffic with some delay requirements to provide a fair allocation of fair transmission media. Although a discussion has assumed the service of Internet protocol packets, it can also be applied with any packet data service which must meet some delay and/or jitter constraints.

[0031] Numerous additional modifications and variations of the present